

REMARKS

Applicant appreciates the time taken by the Examiner to review Applicant's present application. This application has been carefully reviewed in light of the Official Action mailed July 30, 2008. Applicant respectfully requests reconsideration and favorable action in this case.

Interview Summary

Pursuant to Applicant Initiated Interview Request submitted September 19, 2008, a telephonic interview was conducted on October 2, 2008 between Examiners Richard Keehn, Dohm Chankong and Agent Kevin Gust. During the interview, Examiners Keehn and Chankong and Agent Gust discussed the 102 rejection of Claim 1, the 103 rejection of claim 60, the 112 rejection of claim 6 and the 101 rejection of claim 72. The inventive concepts and differences to the prior art were discussed. Applicant appreciates the time and effort taken by Examiners Keehn and Chankong to review Applicant's present application and discuss the pending claims and the cited prior art.

Rejections under 35 U.S.C. § 101

Claim 72 was rejected under 35 U.S.C. § 101. Claim 72 is cancelled herein.

Rejections under 35 U.S.C. § 112

Claim 6 was rejected under 35 U.S.C. § 112, second paragraph as indefinite due to the phrase "provides receives user input." Claim 6 is amended herein. Accordingly, withdrawal of this rejection is requested.

Rejections under 35 U.S.C. § 102

Claims 1-3 and 6-28 were rejected as anticipated by U.S. Patent No. 6,498,786 ("Kirkby"). Applicant respectfully traverses the rejection.

As amended, Claim 1 recites:

A system for allocating resources amongst a plurality of applications, the system comprising:

a distributed server pool director for organizing and maintaining a set of servers in a server pool;

a monitoring module at each server in the server pool for detecting demands for one or more resources located on the server and exchanging information regarding demands for the one or more resources at the plurality of servers;

a distributed policy engine for specifying a policy for allocation of resources of the plurality of computers amongst the plurality of applications having access to the resources, wherein the policy is based on the demands for the one or more resources and adapted based on changes in the behavior of the system; and an enforcement module at each computer for allocating the resources amongst the plurality of applications based on a decision made by the policy engine.

Thus, embodiments of systems for allocating resources amongst a plurality of computers may include a distributed pool director for organizing and maintaining a set of servers in a server pool, a monitoring module for detecting demands for one or more resources located on that server and exchanging information regarding demands for the one or more resources, a distributed policy engine for specifying a policy for allocating resources at the plurality of servers based on the demand for the resources and adapted based on changes to the system, and an enforcement module at each computer for allocating the resources based on a decision made by the policy engine.

Embodiments of the distributed system can receive input from a variety of monitoring sources which describe aspects of the state and performance of applications running in the computing environment as well as the underlying resources (e.g., computer servers) which are servicing the applications. (See, Specification, para. 63.) Applicant respectfully submits that Kirkby does not appear to allocate resources as disclosed by Applicant. Instead, Kirkby describes a system which receives customer input and determines a willingness-to-pay (WTP), such that allocation of resources is made based on a customer's willingness to pay for more resources. (See, Kirkby, Col. 5, lines 4-9.) Furthermore, Applicant respectfully submits that the Kirkby reference fails to teach or describe a policy engine for specifying a policy for the allocation of resources, wherein the policy is based on the demands for the one or more resources and adapted based on changes in the behavior of the system. The methods described by Kirkby allocate resources to each of several "source-sink" entities based on a numerical value (termed "Willingness To Pay, or "WTP") associated with each entity.

For at least the foregoing reasons, Applicant respectfully submits that claim 1 and correspondingly the dependent claims recite subject matter not described by Kirkby under 35 U.S.C. § 102(b) and therefore should be allowed. Accordingly, withdrawal of this rejection is respectfully requested.

Rejections under 35 U.S.C. § 103

Claims 4 and 5 were rejected as obvious over Kirkby in view of U.S. Patent No. 5,414,845 (“Behm”). In the rejection, the Examiner states that Behm discloses an invention substantially as claimed, including wherein the resources include processing resources. Applicant traverses the rejection.

As mentioned above above, Applicant respectfully submits that the Kirkby reference fails to teach or describe a policy engine for specifying a policy for the allocation of resources, wherein the policy is based on the demands for the one or more resources and adapted based on changes in the behavior of the system. Furthermore, the methods described by Kirkby calculate the fractions of a set of shared resources (i.e., bandwidth) to be allocated. In contrast, Applicant claims embodiments which support a broader set of action types to improve quality of service (QoS) to clients, including the allocation of not only bandwidth, but also CPU, memory, servers to distributed applications, power management of servers, the starting and stopping of application instances, and executing application-specific, user-defined commands that are supplied to the system dynamically. (See, Specification, para. 64.)

In the rejection, the Examiner states that Behm teaches wherein the resources include processing resources. Applicant respectfully submits that the teachings of Behm fail to remedy the deficiencies of Kirkby. Behm describes a network-based computer system for scheduling which request Go processes next and on which batch mode. (See, Behm, Abstract, Figure 3.) Applicant respectfully submits that the necessity of waiting for events is one of the shortcomings that may be overcome with a distributed system. Thus, the processing resources used by Behm have a different function (e.g., scheduling) than the processing resources used by Applicant (e.g., allocating resources). For at least these reasons, Behm appears to teach away from allocating resources in a distributed system according to a policy based on demands for resources and applications as recited in the claims. Accordingly, withdrawal of this rejection is requested.

Claims 60-64, 66-67 and 70-72 were rejected as obvious over U.S. Patent No. 6,766,348 (“Combs”) in view of U.S. Publication No. 2003/0069972 (“Yoshimura”). Applicant traverses the rejection and respectfully submits that allocation performed by Combs differs substantially from allocation performed by embodiments disclosed by Applicant.

As amended, claim 60 recites:

A method for allocating resources to a plurality of applications, the method comprising:

receiving user input specifying priorities of the plurality of applications to resources of a plurality of servers, the specified priorities including designated servers assigned to at least some of the plurality of applications;

selecting a given application based upon the specified priorities of the plurality of applications;

determining the demand for one or more resources located on each server in a plurality of communicatively connected servers;

specifying a policy for allocation of resources of the plurality of servers;

allocating to the application resources located on one or more of the plurality of servers based on the policy for allocation of resources;

allocating additional resources to the application until the application's demands for resources are satisfied; and

repeating above steps for each of the plurality of applications based on the specified priorities.

Thus, embodiments for allocating resources to a plurality of applications may include receiving user input specifying priorities of the applications, selecting an application based upon the specified priorities, determining the demand for one or more resources, specifying a policy for allocation of resources, allocating resources to the application resources located on the one or more servers, and allocating additional resources located on the plurality of servers.

In the rejection, the Examiner states that Combs discloses selecting an application based upon the specified priorities of the plurality of applications, in which a Resource Allocator Handling System (RAHS) performs services in order of priority from highest to lowest. However, in Coombs, resource allocation is generally resource-centric, in that the system forms a resource list by sorting the resources according to the greatest remaining capacity. (See, Combs, Col. 13, lines 23-27, Col. 14, lines 52-54.) In contrast, embodiments disclosed by Applicant may be application-centric, in that the system selects a given application based on the specified priorities of the applications. The priorities of the applications may be specified by a user. (See, Specification, para. 73.) Thus, Combs does not specify a policy for the allocation of resources or allocate additional resources to the application until the demands for resources are satisfied.

Applicant respectfully submits that Yoshimura fails to remedy the shortcomings of Combs. Instead of teaching a distributed system that allocates resources such as CPU shares, bandwidth etc., based on priorities of the applications, Yoshimura describes a system in which a user identification table and VPN, VLAN and storage configuration definition tables are managed in a managing server. (See, Yoshimura, para. 11.)

For at least the foregoing reasons, Applicant respectfully submits that Combs and Yoshimura, alone or in combination, fail to teach determining the demand for one or more

resources located on each server in a plurality of communicatively connected servers, specifying a policy for allocation of resources of the plurality of servers, and allocating to the application resources located on one or more of the plurality of servers based on the policy for allocation of resources. Accordingly, withdrawal of this rejection is requested.

Claims 65, 68 and 69 were rejected as obvious over Combs, Yoshimura and further in view of U.S. Publication No. 2005/0177755 ("Fung"). The rejection is traversed.

Notwithstanding the arguments presented above, Applicant respectfully submits that Fung fails to overcome the shortcomings of Combs and Yoshimura and submits that Fung teaches only power management and fails to teach allocating resources. The reasons for shutting down power in the Fung patent may be very different than Applicant's shutting down a resource. For example, Fung describes power designs so that even when the server is operating at its maximum performance level and consuming its maximum power, that consumption may still be lowered. (See, Fung, para. 46.) However, Applicant discloses embodiments in which resources may also include bandwidth, CPU shares, etc. Thus, for example, embodiments of Applicant's invention may decrease the bandwidth used by one application so that another application may have more bandwidth. (See, Specification, para. 64.) For at least the foregoing reasons, Applicant respectfully submits that Fung, alone or in combination with Combs and Yoshimura, fails to teach or describe allocating resources as recited in claims 65, 68 and 69. Accordingly, withdrawal of this rejection is respectfully requested.

Conclusion

Applicant has now made an earnest attempt to place this case in condition for allowance. Other than as explicitly set forth above, this reply does not include an acquiescence to statements, assertions, assumptions, conclusions, or any combination thereof in the Office Action. For the foregoing reasons and for other reasons clearly apparent, Applicant respectfully requests full allowance of Claims 1-28 and 60-71. The Examiner is invited to telephone the undersigned at the number listed below for prompt action in the event any issues remain.

The Director of the U.S. Patent and Trademark Office is hereby authorized to charge any fees or credit any overpayments to Deposit Account No. 50-3183 of Sprinkle IP Law Group.

Respectfully submitted,

Sprinkle IP Law Group
Attorneys for Applicant

Ari G. Akmal
Reg. No. 51,388

Date: October 30, 2008

1301 W. 25th Street, Suite 408
Austin, TX 78705
Tel. (512) 637-9220
Fax. (512) 371-9088